arXiv:1412.8005v1 [cs.DS] 27 Dec 2014

Dynamic Algorithms for Interval Scheduling on a Single Machine

Alexander Gavruskin, Bakhadyr Khoussainov, Mikhail Kokho, Jiamou Liu

1. Introduction

Imagine a number of processes all need to use a particular resource for a period of time. Each process i specifies a starting time $s(i)$ and a finishing time $f(i)$ between which it needs to continuously occupy the resource. The resource cannot be shared by two processes at any instance. One is required to design a scheduler which chooses a subset of these processes so that 1) there is no time conflict between processes in using the resource; and 2) there are as many processes as possible that get chosen.

The above is a typical set-up for the interval scheduling problem, one of the basic problems in the study of algorithms. Formally, given a collection of intervals on the real line specified by their starting and finishing times, the problem asks for a subset of maximal size consisting of pairwise non-overlapping intervals. The interval scheduling problem and its variants appear in a wide range of areas in computer science and applications such as in logistics, telecommunication, and manufacturing. They form an important class of scheduling problems and have been studied under various names and with application-specific constraints [9].

The interval scheduling problem, as stated above, can be solved by a *greedy* scheduler as follows [8]. The scheduler sorts intervals based on their finishing time, and then iteratively selects the interval with the least finishing time that is compatible with the intervals that have already been scheduled. The set of intervals chosen in this manner is guaranteed to have maximal size. This algorithm works in a *static* context in the sense that the set of intervals is given a priori and it is not subject to change.

In a *dynamic* context the instance of the interval scheduling problem is usually changed by a real-time events, and a previously optimal schedule may become not optimal. Examples of such real-time events include job cancelation, arrival of an urgent job, and change in job processing time. To avoid the repetitive work of rerunning the static algorithm every time when the problem instance has changed, there is a demand for efficient *dynamic algorithms* for solving the scheduling problem on the changed instances. In this dynamic context, the set of intervals change through a number of *update operations* such as insertion or

removal. Our goal is to design algorithms that allow us to solve the interval scheduling problem in a dynamic setting.

A natural setting for the problem is a special class of interval sets, which we call *monotonic interval sets*. In a monotonic set no interval is properly contained by another interval. For example, if all processes require the same amount of time to be completed, then the set of intervals is monotonic. Moreover, monotonic interval sets are closely related to proper interval graphs. An *interval graph* is an undirected graph whose nodes are intervals and two nodes are adjacent if the two corresponding intervals overlap. A *proper interval graph* is an interval graph for a monotonic set of intervals. There exist linear time algorithms for representing a proper interval graph by a monotonic set of intervals [1, 6, 2]. Furthermore, solving the interval scheduling problem for monotonic intervals corresponds to finding a maximal independent set in a proper interval graph.

1.1. Related work.

On a somewhat related work, S. Fung, C. Poon and F. Zheng [3] investigated an online version of interval scheduling problem for weighted intervals with equal length (hence, the intervals are monotonic), and designed randomised algorithms. We also mention that R. Lipton and A. Tompkins [5] initiated the study of online version of the interval scheduling problem. In this version a set of intervals are presented to a scheduler in order of start time. Upon seeing each interval the algorithm must decide whether to include the interval into the schedule.

A related problem on a set of intervals I asks to find a minimal set of points S such that every interval from I contains at least one point from S. Such a set S is called a *piercing set* of I. A dynamic algorithm for maintaining a minimal piercing set S is studied in [4]. The dynamic algorithm runs in time $O(|S| \log |I|)$. We remark here that if one has a maximal set J of disjoint intervals in I, one can use J to find a minimal piercing set of I, where each point in the piercing set corresponds to the finishing time of an interval in J in time $O(|J|)$. Therefore our dynamic algorithm can be adapted to one that maintains a minimal piercing set. Our algorithm improves the results in [4] when the interval set I is monotonic.

Kaplan et al. in [7] studied a problem of maintaining a set of nested intervals with priorities. The problem asks for an algorithm that given a point p finds the interval with maximal priority containing p. Similarly to our dynamic algorithm, the solution in [7] also uses dynamic trees to represent a set of intervals.

1.2. Our results.

For the monotonic case, we provide two dynamic algorithms solving the interval scheduling problem. The first algorithm has $O(\log^2 n)$ amortised complexity for update operations and $O(\log n)$ amortised complexity for the query operations. The second algorithm improves the complexity of update operations to $O(\log n)$ amortised. For the general case, we extend the first algorithm. The complexity of the query operation remains the same, while the complexity of the

update operation increases by the factor of d, where d is the maximal number of overlapping intervals. Formal explanation are in the next sections.

The first algorithm maintains the *compatibility forest* data structure denoted by CF. We say the *right compatible interval* of an interval i is the interval j such that $f(i) < s(j)$ and there does not exist an interval ℓ such that $f(i) < s(\ell)$ and $f(\ell) < f(j)$. The CF data structure maintains the right compatible interval relation. The implementation of the data structure utilises, nontrivially, the dynamic tree data structure of Sleator and Tarjan [11]. As a result, in **Theorem 4** and **Theorem 7** we prove the amortised bounds for the monotonic and non-monotonic interval sets respectively.

The second dynamic algorithm maintains the *linearised tree* data structure denoted by LT. We say that intervals are *equivalent* if their right compatible intervals coincide. The LT data structure maintains both the right compatibility relation and the equivalence relation.Then, in **Theorem 13** of Section 4 we prove that the insertion, removal and query operations take time amortised $O(\log n)$. However, this comes with a cost. As opposed to the CF data structure that keeps a representation of an optimal set after each update operation, the linearised tree data structure does not explicitly represent the optimal solution.

To test the performance of our algorithms for the monotonic interval sets, we carried out experiments on random sequences of update and query operations. The experiments show that the two data structures CF and LT perform similarly. The reason for this is that the first dynamic algorithm based on CF reaches the bound of $\log^2 n$ only on specific sequences of operations, while on uniformly random sequences the algorithm may run much faster.

Organisation of the paper. Section 2 introduces the problem and monotonic interval sets. Section 3 describes the compatibility forest data structure and algorithms for monotonic and non-monotonic interval sets. Section 4 describes the linearsed tree data structures and present our second dynamic algorithm, which is based on the linearised tree. Section **??** discusses the experiments.

2. Preliminaries

Interval scheduling basics. An *interval* is a pair $(s(i), f(i)) \in \mathbb{R}^2$ with $s(i) < f(i)$, where $s(i)$ is the *starting time* and $f(i)$ is the *finishing time* of the interval. We abuse notation and write i for the interval $(s(i), f(i))$. Two intervals i and j are *compatible* if $f(i) < s(j)$ or $f(j) < s(i)$. Otherwise, these two intervals *overlap*. Given a collection of intervals $I = \{i_1, i_2, \ldots, i_k\}$, a *compatible set* of I is a subset $J \subseteq I$ such that the intervals in J are pairwise compatible. An *optimal set* of I is a compatible set of maximal size. The *interval scheduling problem* consists of designing an algorithm that finds an optimal set.

We recall the greedy algorithm that solves the problem [8]. The algorithm sorts intervals by their finishing time, and then iteratively chooses the interval with the least finishing time compatible with the last selected interval. The set of thus selected intervals is optimal. The algorithm takes $O(n \log n)$ worst-case time where n is the size of I. If the sorting is already given then the algorithm

runs in linear time. Below, we formally define the greedy optimal set found by this greedy algorithm.

Let $\preceq$ be the ordering of the intervals by their finishing time. Throughout, by the *least interval*, the *greatest interval*, the *next interval*, the *previous interval*, we mean the least, greatest, next and previous interval with respect to $\preceq$. Without loss of generality we may assume that the intervals in I have pairwise distinct finishing times. Given the collection I, we inductively define the set $J = \{i_1, i_2, \ldots\}$, the *greedy optimal set* of I, as follows. The interval i_1 is the least interval in I. The interval i_{k+1} is the least interval compatible with i_k such that $i_k \prec i_{k+1}$. The set J obtained this way is an optimal set [8].

The set I of intervals is called *monotonic* if no interval in I contains another interval. The *right compatible interval* of i, denoted by $\mathsf{rc}(i)$, is the least interval j compatible with i such that $i \prec j$. Similarly, the *left compatible interval* of i, written $\mathsf{lc}(i)$, is the greatest interval j compatible with i such that $j \prec i$.

2.1. Data Structures

Binary Search Tree. A *binary search tree* is a standard data structure that maintains a linearly ordered collection of records. The data structure supports the operations

$$\mathsf{insert}(T, u),\ \mathsf{delete}(T, u),\ \mathsf{find}(T, u),$$
$$\mathsf{predecessor}(u),\ \mathsf{successor}(u),\ \mathsf{maximum}(T),\ \text{ and }\ \mathsf{minimum}(T),$$

where T is a binary search tree and u is an element from the domain. If a binary search tree is balanced, the complexity of all the above operations is $O(\log n)$ where n is the number of elements in the collection. We point out that there are well-known self-balancing binary search tree data structures such as AVL tree and red-black tree.

Splay Tree. A *splay tree* is also a self-balancing binary search tree for storing linearly ordered objects. In addition to the operations for binary search trees, the splay tree data structure also supports the following operations.

- $\mathsf{splay}(u)$: This operation reorganises a splay tree so that u becomes the root.
- $\mathsf{join}(A, B)$: This operation joins two splay trees A, B into one splay tree, where any interval in A is less than any interval in B, into one tree.
- $\mathsf{split}(A, u)$: This operation splits the splay tree A of u into two new splay trees
$$R(u) = \{x \in A \mid u \leq x\} \text{ and } L(u) = \{x \in A \mid x < u\}.$$

All the operations for splay trees take $O(\log n)$ amortised time [12].

Dynamic Trees. A dynamic tree data structure maintains a collection of objects that are stored in a number of rooted trees, viewed as directed graphs with edges pointing from children to parents. The trees can be manipulated using the following operations:

- $\mathsf{link}(v, u)$: If v is the root of a tree and u is a node in another tree, add an edge from v to u and thus "link" the trees containing v and u together.
- $\mathsf{cut}(v)$: If v is not the root of a tree, delete the edge from v to its parent and thus divide the tree containing v into two.

These operations have $O(\log n)$ amortised time complexity [11].

2.2. Problem Setup

In this setting the collection I of intervals changes over time. Thus, the input to the problem is an arbitrary sequence $o_1, \ldots, o_m$ of update and query operations described as follows:

- *Update operations*: $\mathsf{insert}(i)$ inserts an interval i and $\mathsf{remove}(i)$ removes an interval i.
- *Query operation*: The operation $\mathsf{query}(i)$ returns true if i belongs to the greedy optimal set and false otherwise.

Our goal is to design algorithms for performing these operations that minimise the total running time.

3. Compatibility Forest Data Structure

In this section we define compatibility forest and describe how to maintain efficiently maintain it for a set of intervals. We first show how to represent a monotonic set, and then we extend the algorithms for the general case.

3.1. Definition of Compatibility Forest

Let I be a set of intervals. We define the *compatibility forest* as a graph $\mathcal{F}(I) = (V, E)$ where $V = I$ and $(i, j) \in E$ if $j = \mathsf{rc}(i)$. By a forest we mean a directed graph where the edge set contains links from nodes to their parents. We denote the parent of a node v by $p(v)$. The *roots* and *leaves* are standard notions that we do not define. Figure 1 shows an example of a monotonic set of intervals with its compatibility forest. We note that for every forest one can construct in a linear time a monotonic set of intervals whose compatibility forest coincides (up to isomorphism) with the forest.

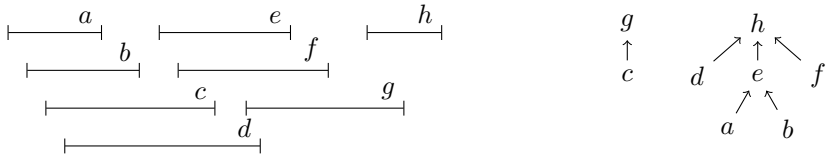

Figure 1: Example of a monotonic set of intervals and its compatibility forest.

A *path* in the compatibility forest $\mathcal{F}(I)$ is a sequence of nodes $i_1, i_2, \ldots, i_k$ where $(i_t, i_{t+1}) \in E$ for any $t = 1, \ldots, k-1$. It is clear that any path in the forest $\mathcal{F}(I)$ consists of compatible intervals. Essentially, the forest $\mathcal{F}(I)$ connects nodes by the greedy rule: for any node i in the forest $\mathcal{F}(I)$, if the greedy rule is applied to i, then the rule selects the parent j of i in the forest. Hence, the longest paths in the compatibility forest correspond to an optimal sets of I. In particular, the path starting from the least interval is the greedy optimal set. Our first dynamic algorithm amounts to maintaining this path in the forest $\mathcal{F}(I)$.

We explain how we maintain paths in the compatibility forest $\mathcal{F}(I)$. The representation of the forest is developed from the dynamic tree data structure as in [11]. The idea is to partition the compatibility forest into a set of node-disjoint paths. Paths are defined by two types of edges, *solid edges* and *dashed edges*. Each node in the compatibility forest is required to have at most one incoming solid edge. A sequence of edges $(u_0, u_1), (u_1, u_2), \ldots, (u_{k-1}, u_k)$ where each (u_i, u_{i+1}) is a solid edge is called a *solid path*. A solid path is *maximal* if it is not properly contained in any other solid path. Therefore, the solid edges in $\mathcal{F}(I)$ form several maximal solid paths in the forest. Furthermore, the data structure ensures that each node belongs to some maximal solid path. There is an important subroutine in the dynamic tree data structure called the *expose* operation [11]. The operation starts from a node v and traverses the path from v to the root: while traversing, if the edge $(x, p(x))$ is dashed, we declare $(x, p(x))$ solid and declare the incoming solid edge (if it exists) incident to $p(x)$ dashed. Thus, after exposing node v, all the edges on the path from v to the root become solid. Note that in CF data structure the $p(x)$ and $\mathsf{rc}(x)$ are the same.

3.2. Compatibility forest of a monotonic set of intervals

We denote the representation of $\mathcal{F}(I)$ for the monotonic interval set by MonCF. The representation consists of two components. The first is a binary search tree $T(I)$. The nodes of $T(I)$ are intervals in I ordered by their starting time. Note that monotonicity of I implies that the order of intervals in $T(I)$ coincide with $\preceq$, order of intervals by their finishing time. In addition to standard operations of binary search trees, we define the right_compatible operation. Given an interval, i the operation returns $\mathsf{rc}(i)$, if it is in I, or nil, otherwise. The second component is a set of splay trees. Each splay tree stores the nodes of a maximal solid path in the compatibility forest $\mathcal{F}(I)$ We denote by ST_i the splay tree containing the interval i.

Algorithm 1 right_compatible(i)

```
1: r ← nil
2: j ← the root in the interval tree T(I).
3: while j ≠ nil do
4:     if j ⪯ i or j overlaps i then
5:         j ← the right child of j
6:     else
7:         r ← j
8:         j ← the left child of j
9: return r
```

Lemma 1. *On monotonic set I of intervals the operation* right_compatible(i) *run in time* $\Theta(\log n)$ *and return* rc(i).

To prove the lemma we observe that for a monotonic set I of intervals and $i, j \in I$, if i overlaps j, then each of the intervals between i and j overlaps both i and j.

Proof. For the complexity, note that the length of paths from a leaf to the root in $T(I)$ is $\lfloor \log n \rfloor + 1$. Thus, the operation takes time $\Theta(\log n)$.

For the correctness, we use the following loop invariant: *If I contains* rc(i)*, then the subtree rooted at j contains* rc(i) *or r equals* rc(i).

Initially, j is the root of $T(I)$, so the invariant holds. Each iteration of the **while** loop executes either line 5 or lines 7-8 of Alg. 1. If line 5 is executed, then we have $j \preceq i$ or j overlaps i. If $j \preceq i$ then all intervals in the left subtree of j are less than i. If $j \succeq i$ but j overlaps i, then by the observation above, all intervals between i and j overlap i. In both cases, none of the intervals in the left subtree of j is rc(i). Therefore setting j to be the right child of j preserves the invariant.

If lines 7-8 are executed, then we have $j \succeq i$ and j is compatible with i. If there exists an interval that is less than j and compatible with i, then such an interval is in the left subtree of j. If such an interval does not exist, j is the smallest interval which is compatible with i. Therefore setting r to be j and j to be the right child of j preserves the invariant.

Thus, the algorithm outputs rc(i) if it exists and outputs nil otherwise. Indeed, the loop terminates when $j =$ nil. Hence if the set of intervals I contains rc(i) then $r =$ rc(i). If I does not contain rc(i) then line 5 is executed at every iteration, so $r =$ nil. □

We now describe algorithms for maintaining compatibility forest data structure. We call the algorithms queryMonCF, insertMonCF and removeMonCF for the query, insertion, and removal operations, respectively.

The operation queryMonCF*:* To perform this operation on an interval i, we first find in the interval tree $T(I)$ the minimum element m. We then check if i belongs to the splay tree ST_m. We return true if $i \in \mathrm{ST}_m$; otherwise we return false.

Algorithm 2 queryMonCF(i)

```
1: m ← minimum(T(I))
2: return find(ST_m, i) = i
```

The operation expose*:* To expose an interval i, we find the maximum element j in the splay tree ST_i. Then find the right compatible interval $i' = \mathsf{rc}(j)$. If i' does not exist (that is, j is a root in the compatibility forest), we stop the process. Otherwise, (j, i') is a dashed edge. We split the splay tree at i' into trees $L(i')$ and $R(i')$ and join ST_i with $R(j')$. We then repeat the process taking i' as i.

Algorithm 3 expose(i)

```
1: j ← maximum(ST_i)
2: i' ← right_compatible(j)
3: while i' is not nil do
4:     split(ST_{i'}, i')
5:     join(ST_i, R(i'))
6:     j ← maximum(ST_{i'})
7:     i' ← right_compatible(j)
```

The operation insertMonCF*:* To insert an interval i, we add i into the tree $T(I)$. Then we locate the next interval r of i in the ordering $\preceq$. If such r exists, we access r in the splay tree ST_r and find the interval j such that (j, r) is a solid edge. If such a j exists and j is compatible with i, we delete the edge (j, r) and create a new edge (j, i) and declare it solid. We restore the longest path of the compatibility forest by exposing the least interval in $T(I)$.

Algorithm 4 insertMonCF(i)

```
1: insert(T(I), i)
2: r ← next(i)                                  ▷ Find the next interval of i
3: if r ≠ nil then
4:     j ← predecessor(ST_r, r)                 ▷ Find a solid edge (j, r)
5:     if j ≠ nil and j is compatible with i then
6:         split(ST_r, r)                       ▷ Destroy the solid edge (j, r)
7: expose(mininum(T(I)))
```

The operation removeMonCF*:* To delete an interval i, we delete the incoming and outgoing solid edges of i if such edges exist. We then delete i from the tree $T(I)$. We restore the longest path of the CF by exposing the least interval in $T(I)$.

Algorithm 5 removeMonCF(i)

1: remove(ST_i, i) ▷ Delete i from its splay tree ST_i
2: remove($T(I), i$) ▷ Delete i from the interval tree $T(I)$
3: expose(minimum($T(I)$))

Correctness. For correctness of operations, we use the following invariants.

(A1) Every splay tree represents a maximal path formed from solid edges.

(A2) Let m be the least interval in I. The splay tree ST_m contains all intervals on the path from m to the root.

Note that (A2) guarantees that the query operation correctly determines if a given interval i is in the greedy optimal set. The next lemma shows that (A1) and (A2) are invariants indeed and that the operations correctly solve the dynamic monotonic interval scheduling problem.

Lemma 2. *(A1) and (A2) are invariants of* insertMonCF*,* removeMonCF*, and* queryMonCF*.*

Proof. For (A1), first consider the operation of joining two splay trees A and B via the operation expose(i). Let j be the maximal element in A and j' be the minimum element in B. In this case, j' is obtained by the operation right_compatible(j). It is clear that (j, j') is an edge in the forest $\mathcal{F}(I)$. Next, consider the case when we apply insertMonCF(i) into the splay tree A. In this case, A is $L(r)$ where r is the next interval of i in I. Let j be the previous interval of r in the tree ST_r. By (A1), before inserting i, (j, r) is an edge in $\mathcal{F}(I)$ and thus $r = \mathsf{rc}(j)$. Note we only insert i to $L(r)$ when j is compatible with i. Since $i < r$, after inserting i, i becomes the new right compatible interval of j. So, joining $L(r)$ with i preserves (A1). Operations removeMonCF(i) and queryMonCF(i) do not create new edges in splay trees. Thus, (A1) is preserved under all operations.

For (A2), the expose(i) operation terminates when it reaches a root of the compatibility forest. As a result, ST_i contains all nodes on the path from i to the root. Since expose(minimum($T(I)$)) is called at the end of both insertMonCF(i) and removeMonCF(i) operations, (A2) is preserved under every operation. □

Complexity. Let n be the number of intervals in I. As discussed in Section 2, all operations for the interval tree have $O(\log n)$ worst case complexity, and all operations for splay trees have $O(\log n)$ amortised complexity. The query operation, involves finding the minimum interval in $T(I)$ and searching i in a splay tree. Hence, the query operation runs in amortised time $O(\log n)$. For each insert and remove operation, we perform a constant number of operations on $T(I)$ and the splay trees plus one expose operation.

To analyse expose operation, define the size size(i) of an interval i to be the number of nodes in the subtree rooted at i in $\mathcal{F}(I)$. Call an edge (i, j) in $\mathcal{F}(I)$ *heavy* if $2 \cdot \mathsf{size}(i) > \mathsf{size}(j)$, and *light* otherwise. It is not hard to see that this partition of edges has the following properties:

($\star$) Every node has at most one incoming heavy edge.

($\star\star$) Every path in the compatibility forest consists of at most $\log n$ light edges.

Lemma 3. *In a sequence of k update operations, the total number of dashed edges, traversed by* expose *operation, is $O(k \log n)$.*

Proof. The number of iterations in expose operation is the number of dashed edges in a path from the least interval to the root. A dashed edge is either heavy or light. From ($\star\star$), there are at most $\log n$ light dashed edges in the path. To count the number of heavy dashed edges, consider the previous update operations. After deletion of i, all children of i become children of the next interval of i. After inserting i, the children of the next interval of i that are compatible with i become children of i. Therefore , there are at most two path where an update operation transforms light dashed edges to heavy dashed edges. Figure 2 illustrates these structural changes. Since there are at most $\log n$ light dashed edges on each path, an update operations creates at most $\log n$ heavy dashed edges.

Execution of expose in an update operation creates at most $\log n$ heavy dashed edges from heavy solid edges. Hence, the total number of heavy dashed edges created after k update operations is $O(k \log n)$. □

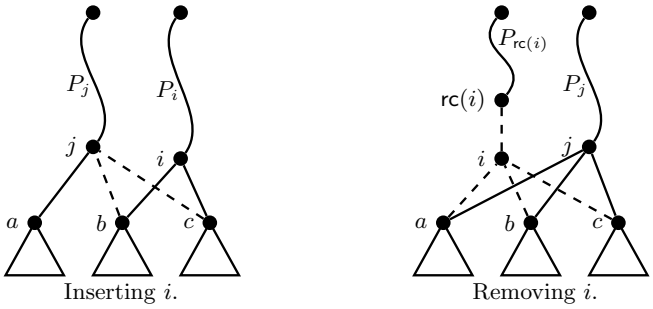

Figure 2: Redirections of edges in CF, where j is the next interval of i.

Lemma 2 and Lemma 3 give us the following theorem:

Theorem 4. *The algorithms* queryMonCF, insertMonCF *and* removeMonCF *solve the dynamic monotonic interval scheduling problem. The algorithms perform insert interval and remove interval operations in $O(\log^2 n)$ amortised time and query operation in $O(\log n)$ amortised time, where n is the size of the set I of intervals.*

Remark. Tarjan and Sleator's dynamic tree data structure has amortised time $O(\log n)$ for update and query operations. To achieve this, the algorithm maintains dashed edges explicitly. Their technique cannot be adapted directly to CF

because insertion or removal of intervals may result in redirections of a linear number of edges. An example is depicted on the Figure 2.Therefore, more care should be taken; for instance, one needs to maintain dashed edges implicitly in $T(I)$ and compute them calling right_compatible operation.

Proposition 5 (Sharpness of the $\log^2 n$ bound). *In* CF *data structure there exists a sequence of k update operations with $\Theta(k \log^2 n)$ total running time.*

Proof. Consider a sequence which creates a set of $n < k$ intervals. We assume that $n = 2^{h+1} - 1$, where $h \in \mathbb{N}$. The first n operations of the sequence are insertMonCF such that the resulted compatibility forest is a perfect binary tree T_n, that is, each internal node of T_n has exactly two children and the height of each leaf in T_n is h. The next $k - n$ operations starting form T_n are pairs of insertMonCF followed by removeMonCF. At stage $s = n + 2m + 1$, insertMonCF inserts an interval i_s into T_s producing the tree T_{s+1}. The interval i_s is such that in T_{s+1} the path from i_s to the root is of length $h+1$ and the path consists of dashed edges only. Then, at stage $s + 1$ we delete i_s. This produces a tree T_{s+2} which is a perfect binary tree of height h. We repeat this $k - n$ times. We can select i_s as desired since each perfect binary tree T_s always has a path of length h consisting of dashed edges only. Therefore a sequence of k such operations takes time $\Theta(k \log^2 n)$. □

3.3. Compatibility forest of a non-monotonic set of intervals

In this section, we show how to maintain compatibility forest for a set of non-monotonic intervals. In this case, extra care should be taken when we insert interval i that is covered by other intervals since i may become the new right compatible interval for several overlapping intervals. The example in Figure 3 shows such insertion. Therefore, when we insert an interval i, we need to find all intervals covering i.

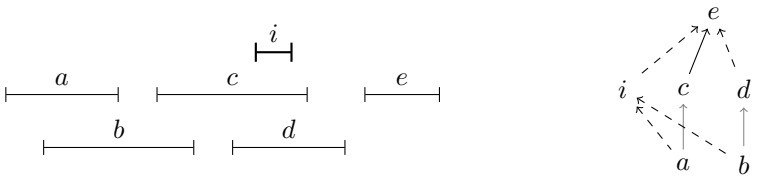

Figure 3: Insertion of i into CF destroys solid edges (a, c) and (b, d)

We describe a new operation covers(i), which returns all intervals covering a given interval. To support this operation, we introduce the third component - an interval tree data structure. An *interval tree* [10] is a leaf-oriented balanced binary search tree where leaves store endpoints of the intervals in increasing order. Intervals themselves are stored in the internal nodes as follows. For each internal node v the set $I(v)$ consists of intervals that contain the *split point* of v

and are covered by the *range* of v. The split point of v, denoted by $split(v)$, is a number such that the leaves of the left subtree of v store endpoints smaller than $split(v)$, and the leaves of the right subtree of v store endpoints greater than $split(v)$. The range of v, denoted by $range(v)$, is defined recursively as follows. The range of the root is $(-\infty, \infty]$. For a node v, where $range(v) = (l, r]$, the range of the left child of v is $(l, split(v)]$, and the range of the right child of v is $(split(v), r]$.

To allow insertions and deletions of intervals in the interval tree, we represent it as a red-black tree IT. In a red-black tree, insertion or deletion of a node takes $O(\log n)$ time plus the time for at most 3 rotations to restore the balance. When performing a rotation around an edge $(v, p(v))$ the sets $I(v)$ and $I(p(v))$ change. Let the range of $p(v)$ be $(\ell, r]$. If v is the left child, the range of $p(v)$ after rotation becomes $[split(v), r]$. If v is the right child, the range of p shortens at the other end and becomes $[\ell, split(v)]$. Therefore all intervals in $I(p(v))$ that intersects with $split(v)$ must be moved to $I(v)$. Note that ranges of other nodes are not affected. We represent $I(v)$ with two binary search trees. The first tree $T_s(v)$ stores intervals of $I(v)$ sorted from left to right by their starting point. The second tree $T_f(v)$ stores intervals of $I(v)$ sorted from right to left by their finishing point. To move intervals from $I(v)$ to $I(u)$, we perform join and split operations on the trees representing these sets. Thus in total we need $O(\log n)$ time to insert or delete a node in IT.

To find all intervals covering interval i we do the following. We walk down in the interval tree starting at the root. At every node v, we compare $split(v)$ with $s(i)$ and $f(i)$. If $split(v) > f(i)$ or $split(v) < s(i)$, we respectively traverse the tree $T_s(v)$ or $T_f(v)$ from left to right and report all intervals that covers i. We continue to the left or right child of v respectively. Otherwise, the split point $split(v)$ intersect with i. We traverse $T_s(v)$, report intervals covering i. We terminate the search at this node, because i intersects with ranges of both children of v, namely $[\ell, split(v))$ and $[split(v), r)$, and therefore cannot be covered by any interval, fully contained in these ranges. Formally, the operation is described in Algorithm 6.

Algorithm 6 covers(i)

```
1: S ← nil
2: v ← root of IT
3: while v ≠ nil do
4:     if split(v) < s(i) then
5:         x ← minimum(T_f(v))
6:         v ← right child of v
7:     else
8:         x ← minimum(T_s(v))
9:         v ← left child of v
10:    while x covers i do
11:        Add i into S
12:        x ← successor(x)
13:    if split(v) ∈ i then
14:        break
15: return S
```

Another thing we need to take care of is the right_compatible operation. Since the set of intervals is not monotonic, the observation, essential for the proof of Lemma 1, does not hold. Namely, there might exists intervals $i \preceq j$ and k covered by j such that i intersects with j, but i is compatible with k. To overcome this difficulty, we augment the search tree $T(I)$ and the operation right_compatible as follows. In every node v of $T(I)$ we keep a pointer to the interval m in the subtree rooted at v with the smallest finishing time. It is not hard to check that these pointers can be updated after a rotation in constant time. Therefore, maintenance of these pointers does not change the asymptotic complexity of operations on $T(I)$.

Recall that when we search for the right compatible interval of i in the monotonic set, we go to the left child of the current interval j if $i \preceq j$ and i is compatible with j. In a non-monotonic set, we need to check if there is an interval k that is covered by j and hence compatible with i. If such k exists, it is in the right subtree of j and we can access it in constant time using the pointers we described above. Therefore, if we go to the left child of x, we remember an interval with the smallest finishing time among three intervals: the last remembered interval, an interval at x or an interval with the smallest finishing time in the right subtree of x. For example, if we search for rc(a) in the tree shown in Figure 4, we traverse the path $\{g, d, b, c\}$ and remember intervals h, d, d one after another.

We are now ready to describe the operations, that maintain a compatibility forest for a non-monotonic set of intervals. The query operation queryCF(i) and the remove operation removeCF(i) are identical to queryMonCF(i) and removeMonCF(i) respectively. The insert operation insertCF(i) does the following. First, we add i into the trees $T(I)$ and IT. Second, as in the monotonic case, we check is there exists a solid edge (j, r) such that i substitutes r. Namely,

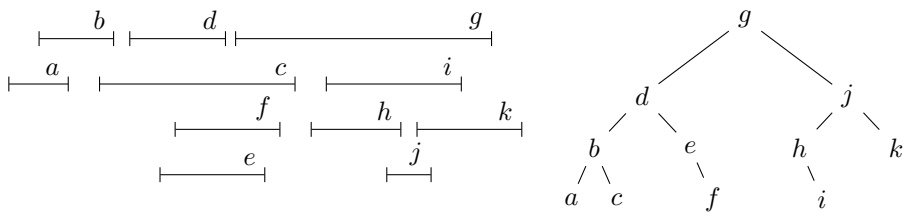

Figure 4: Figure

we search for r such that: (i) $i \prec r$, (ii) i is not covered by r, (iii) for every $i \prec \ell \prec r$, ℓ covers i. Then, if there exist a solid edge from j to r and j is compatible to i, we make this edge dashed. Third, for every interval, which covers i, we make the incoming solid edge, if any, dashed. Finally, we restore the longest path of the compatibility forest by exposing the $\preceq$-least interval. Algorithm 7 describes the operation in details.

Algorithm 7 insertCF(i)

```
1: insert(T(I), i)
2: insert(IT, i)
3: r ← ⪯-next interval that does not cover i
4: if r ≠ nil then
5:     j ← predecessor(ST_r, r)                ▷ Find a solid edge (j, r)
6:     if j ≠ nil and j is compatible with i then
7:         split(ST_r, r)
8: for c in covers(i) do
9:     j ← predecessor(ST_c, c)
10:    if j ≠ nil and j is compatible with i then
11:        split(ST_c, c)
12: m ← ⪯-least interval
13: expose(m)
```

Lemma 6. *In a sequence of k update operations, the total number of dashed edges, traversed by* expose *operation, is $O(d \cdot k \log n)$, where d is the size of a maximal subset of pairwise overlapping intervals.*

Proof. The proof is similar to the proof of Lemma 3. Recall, that an edge (i, j) in the compatibility forest is heavy if the number of nodes in the tree rooted at i is two times greater than the number of nodes in a tree rooted at j. We count the number of heavy dashed edges created by the sequence of k update operations.

Let I_t be the interval after performing t operations. We define d as follows

$$d = \max_{1 \le t \le k} \{|J| \mid J \subseteq I_t \text{ and for every } i, j \in J,\ i \text{ overlap } j\}.$$

Let i be deleted or inserted interval. In the monotonic case, there is at most one interval that exchanges children with i. Here we need to take into account all the intervals $j_1, \ldots, j_m$, that covers i. When we insert i, all children of each j_t becomes children of i. When we delete i, the children of i are distributed among the intervals, covering i. Since there are at most d such intervals, and every path from j_t to the root has at most $\log n$ light dashed edges, an update operation creates at most $d \cdot \log n$ heavy dashed edges. Thus, in total expose operation traverses $O(d \cdot k \log n)$ dashed edges in a sequence of k operations. □

Theorem 7. *The algorithms* queryCF, insertCF *and* removeCF *solve the dynamic interval scheduling problem. The algorithms perform insert interval and remove interval operations in $O(d \log^2 n)$ amortised time and query operation in $O(\log n)$ amortised time, where n is the size of the set I of intervals and d is the size of a maximal subset of pairwise overlapping intervals.*

4. Linearised Tree Data Structure

In this section, we develop a new data structure for the dynamic interval scheduling problem. The dynamic algorithm based on this data structure performs all operations in amortised $O(\log n)$ time. However, the algorithm requires the interval set to be monotonic at all times.

4.1. Definition of Linearised Tree

We say that intervals i and j are *equivalent*, written as $i \sim j$, iff $\mathsf{rc}(i) = \mathsf{rc}(j)$. Denote the equivalence class of i by $[i]$. Thus, two intervals are in the same equivalence class if they are siblings in the compatibility forest. In the linearised tree we arrange all intervals in an equivalence class in a path using the $\preceq$-order. The linearised tree consists of all such "linearised" equivalence classes joined by edges. Hence, there are two types of edges in the linearised tree. The first type connects intervals in the same equivalence class. The second type joins the greatest interval in an equivalence class with its right compatible interval. Formally, the *linearised tree* $\mathcal{L}(I)$ is a triple $(I; E_\sim, E_c)$, where $E_\sim$ and E_c are disjoint set of edges such that:

- $(i, j) \in E_\sim$ if and only if $i \sim j$ and i is the previous interval of j. Call i the *equivalent child* of j.
- $(i, j) \in E_c$ if and only if i is the greatest interval in $[i]$ and $j = \mathsf{rc}(i)$. Call i the *compatible child* of j.

Figure 5 shows an example of a linearised tree. We stress three crucial differences between the CF and LT data structures. The first is that a path in a linearised tree may not be a compatible set of intervals. The second is that linearised trees are binary. The third is when we insert or remove an interval we need to redirect at most two existing edges in the linearised tree. We explain the last fact in more details below when we introduce the dynamic algorithm.

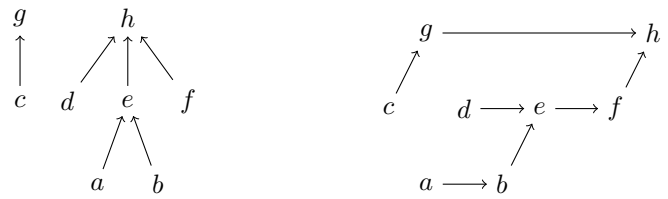

Figure 5: Example of a compatibility forest (left) and linearised tree (right).

We use the dynamic tree data structure to represent the linearised tree. We also maintain the interval tree $T(I)$ as an auxiliary data structure. The interval tree is used to compute previous and next intervals as well as left compatible and right compatible intervals of a given interval.

4.2. Maintaining Linearised Tree

To maintain LT we will need $\mathsf{left_compatible}(i)$ operations, which returns $\mathsf{lc}(i)$ if it is in the interval set or nil otherwise. Algorithms 8 defines this operation.

Algorithm 8 $\mathsf{left_compatible}(i)$

```
1: ℓ ← nil
2: j ← the root in the interval tree T(I).
3: while j ≠ nil do
4:     if j ⪰ i or j overlaps i then
5:         j ← the left child of j
6:     else
7:         ℓ ← j
8:         j ← the right child of j
9: return ℓ
```

We now describe algorithms for maintaining linearised tree data structure. We call the algorithms $\mathsf{queryLT}$, $\mathsf{insertLT}$ and $\mathsf{removeLT}$ for the query, insertion, and removal operations, respectively.

The operation $\mathsf{queryLT}$*:* To detect if an interval i is in the greedy optimal set, consider the path P from the least node m to the root in the linearised tree $\mathcal{L}(I)$. If $i \notin P$, return false. Otherwise, consider the direct predecessor j of i in the path P. If j does not exist or $(j, i) \in E_c$, return true. Otherwise, we return false.

Algorithm 9 queryLT(i)

```
1: m ← minimum(T(I))
2: if i = m then                                  ▷ i is the least interval
3:     return true
4: expose(m)                          ▷ Make the path from m to the root solid
5: if i ≠ find(ST_m, i) then          ▷ i is not on the path from m to the root
6:     return false
7: j ← predecessor(ST_m, i)                       ▷ (j, i) is an edge in LT
8: if i is compatible with j then
9:     return true
10: else
11:     return false
```

Lemma 8. *The operation* queryLT(i) *returns* true *if and only if a given interval* i *belongs to the greedy optimal set of* I*.*

Proof. Let J be the greedy optimal set of I and m be the least element of I. Suppose the algorithm queryLT(i) outputs true. This can happen when (1) $i = m$. In this case i is the least element of I, hence i belongs to J; or (2) i is compatible with predecessor(ST_m, i). Note that for every interval x from ST_m there exists an interval y from the greedy set J such that $y \in [x]$. Consider such an interval $y \in J \cap [\mathsf{predecessor}(\mathrm{ST}_m, i)]$. Since $y \in J$ and i is the next compatible interval of y, i belongs to the greedy optimal set J.

It is not hard to see by induction on the number of elements in J that $J \subseteq \mathrm{ST}_m$. Suppose the algorithm queryLT(i) outputs false. It happens in two cases. First, i is not in ST_m. Then $i \notin J$. Second, predecessor(ST_m, i) exists and is not compatible with i. Then i is not the least interval in $[i] \cap \mathrm{ST}_m$, but every element $x \in J$ is the least element in $[x] \cap \mathrm{ST}_m$. Hence $i \notin J$. □

The operation insertLT*:* Given i, we insert i into $T(I)$. If i is the greatest interval in $[i]$, then we add the edge $(i, \mathsf{rc}(i))$ into E_c. Otherwise, we add the edge (i, j) to $E_\sim$, where j is the next interval equivalent to i. If i has an equivalent child k then we add the edge (k, i) to $E_\sim$ and delete the old outgoing edge from k in case such edge exists. If i has a compatible child ℓ then we add the edge (ℓ, i) to E_c and delete the old outgoing edge in case such edge exists.

Algorithm 10 insertLT(i)

```
1: insert(T(I), i)
2: if i is not the greatest interval in I ∪ {i} then          ▷ i has a parent
3:     if lc(rc(i)) = i then                                  ▷ (i, rc(i)) ∈ E_c
4:         link(i, rc(i))
5:     else                                                   ▷ (i, next(i)) ∈ E_∼
6:         link(i, next(i))
7: j ← previous(i)
8: if rc(j) = rc(i) then                                      ▷ (j, i) ∈ E_∼
9:     cut(j) and link(j, i)
10: j ← lc(i)
11: if rc(j) = i then                                         ▷ (j, i) ∈ E_c
12:     cut(j) and link(j, i)
```

The operation removeLT*:* Given i, we delete i from $T(I)$. We delete an edge from i to the parent of i and redirect the edge from the equivalent child j of i to the parent of i. Then we redirect an edge from the compatible child ℓ of i. Removing i may add new intervals to the equivalence class of ℓ. Therefore if ℓ is still the greatest interval in the updated equivalence class, we add an edge $(\ell, \mathsf{rc}(\ell)$ to E_c. Otherwise, we add the edge (i, j) to $E_\sim$, where j is the next interval of ℓ.

Algorithm 11 removeLT(i)

```
1: if i is not the root then
2:     cut(i)
3: j ← previous(i)
4: if rc(j) = rc(i) then                                   ▷ (j, i) ∈ E_∼
5:     cut(j)
6:     if i = lc(rc(i)) then                               ▷ rc(i) is a new parent of j
7:         link(j, rc(i))
8:     else if i is not the the root then                  ▷ next(i) is a new parent of j
9:         link(j, next(i))
10: j ← lc(i)
11: if i = rc(j) then                                      ▷ (j, i) ∈ E_c
12:     cut(j)
13:     remove(T(I), i)
14:     k ← next(j)
15:     if j is not the root then
16:         if rc(k) = rc(j) then                          ▷ rc(j) ≠ i as we removed i from T(I).
17:             link(j, k)
18:         else
19:             link(j, rc(j))
20: else
21:     remove(T(I), i)
```

4.3. Correctness of the update operations

To prove correctness of the algorithms above, we state two claims about linearised trees. The first claim allows us to check if the given interval the greatest in its equivalent class. The second claim says that changes of the linearised tree after insertion or deletion of an interval i are local with respect to i. We abuse notation and write $(I; E)$ instead of $(I; E_c, E_\sim)$ and $(I'; E')$ instead of $(I'; E'_c, E'_\sim)$. Which edges are used will be clear from the context.

Claim 9. *An interval i is the greatest in $[i]$ if and only if* $\mathsf{lc}(\mathsf{rc}(i)) = i$.

Proof. Let i be the greatest interval in $[i]$. Then for any $j \in [i]$ we have that $j \preceq i$. Assume that $k = \mathsf{lc}(\mathsf{rc}(i)) \neq i$. Then $i \preceq k$ which is a contradiction.

For the other direction, assume that i is not the greatest interval in its equivalent class, that is there exists $j \in [i]$ such that $i \prec j$. Clearly, j is compatible with $\mathsf{rc}(i)$. Therefore $\mathsf{lc}(\mathsf{rc}(i)) = j$, witch is a contradiction. □

Claim 10. *Let $\mathcal{L}(I) = (I, E)$ and $\mathcal{L}(I') = (I', E')$ be two linearised trees such that $I' = I \cup \{i\}$. Let j and k be intervals from the set I'. Then the following properties are satisfied:*

(1) if $(j, k) \notin E$ and $(j, k) \in E'$, then either $j = i$ or $k = i$.

(2) if $(j, k) \in E$ and $(j, k) \notin E'$, then $(j, i) \in E'$.

PROOF. For the first property, we note that if two intervals from I are not connected by an edge in $\mathcal{L}(I)$ then they are not connected by an edge in a bigger linearised tree $\mathcal{L}(I')$. Hence either $j = i$ or $k = i$. For the second property, let $\ell \neq k$ be a parent of j in $\mathcal{L}(I')$. Because $(j, \ell) \notin E$ and $(j, \ell) \in E'$, the property (1) implies that either $j = i$ or $\ell = i$. Thus $\ell = i$. □

Lemma 11. *The operation* $\mathsf{insertLT}(i)$ *preserves linearised tree data structure.*

PROOF. Consider intervals $j, k \in I'$, where $j \preceq k$ and $I' = I \cup \{i\}$. Let (I', E') be the resulting tree after the algorithm $\mathsf{insertLT}(i)$ is performed. We show that $(j, k) \in E'$ if and only if (j, k) is an edge in $\mathcal{L}(I')$.

($\rightarrow$) Suppose that $(j, k) \in E'$. We prove that (j, k) is an edge in $\mathcal{L}(I')$.

- Let $(j, k) \notin E$. Then the algorithm $\mathsf{insertLT}$ must have added (j, k) into E'. Any edge the algorithm adds is adjacent to i. First, we consider outgoing edges, that is, we consider the case when $j = i$. If the algorithm adds an edge from i to $\mathsf{rc}(i)$, then $i = \mathsf{lc}(\mathsf{rc}(i))$ (see lines 3-4 of the Algorithm 10). By Sublemma 9, i is the greatest interval in its equivalence class. If the algorithm adds an edge from i to the next interval k of i, then i is not the greatest interval in $[i]$ and $k \sim i$ (see lines 3-6). Second, we consider incoming edges, that is, we consider the case when $k = i$. If the algorithm adds an edge from $\mathsf{lc}(i)$ to i, then $j = \mathsf{lc}(\mathsf{rc}(j))$ (see lines 10-12). By Sublemma 9, j is the greatest interval in its equivalence class. If the algorithm adds an edge from the previous interval j of i to i, then $j \sim i$ (see lines 7-9). Note that any of the edges added by the algorithm is an edge in $\mathcal{L}(I')$. Hence (j, k) is an edge in $\mathcal{L}(I')$.

- Let $(j, k) \in E$. Assume that (j, k) is not an edge in $\mathcal{L}(I')$. By Sublemma 10 (j, i) is an edge in $\mathcal{L}(I')$. If j is the equivalent child of i, then j is the previous interval of i and $\mathsf{rc}(j) = \mathsf{rc}(i)$. If j is the compatible child of i, then $j = \mathsf{lc}(\mathsf{rc}(j))$. In both of these cases the algorithm deletes the edge from j (see lines 7-9 and 10-12 correspondently). Thus $(j, k) \notin E'$, which is a contradiction.

($\leftarrow$) Suppose that (j, k) is an edge in $\mathcal{L}(I')$. We prove that $(j, k) \in E'$.

- Let $(j, k) \in E$. Assume that $(j, k) \notin E'$. Then the algorithm $\mathsf{insertLT}$ must have deleted (j, k). There are two cases: j is the previous interval of i and $j \sim i$ (see lines 7-9), or $i = \mathsf{rc}(j)$ and j is the greatest interval in $[j]$ (see lines 10-12). In either case, j is a child of i in $\mathcal{L}(I')$, that is $k = i$, which is a contradiction to the assumption that $(j, k) \in E$.

- Let $(j, k) \notin E$. By Sublemma 10, either $j = i$ or $k = i$. Suppose $j = i$. If i is the compatible child of k in $\mathcal{L}(I')$, then $k = \mathsf{rc}(i)$ and, by Sublemma 9, $i = \mathsf{lc}(\mathsf{rc}(i))$. If i is the equivalent child of k, then k is the next interval i and $k \sim i$. The algorithm $\mathsf{insertLT}$ adds the edge (i, k) to E' in lines 3-6. Suppose, $k = i$. If j is the equivalent child of i, j is the previous interval of i and $j \sim i$. If j is the compatible child of i, then $j = \mathsf{lc}(\mathsf{rc}(j))$. The

algorithm insertLT adds the edges (j, i) to E' in lines 7-12. In any case, the edge $(j, k) \in E'$. □

Lemma 12. *The operation* removeLT(i) *preserves linearised tree data structure.*

PROOF. Suppose $\mathcal{L}(I) = (I, E)$ is the linearised tree of a set I of intervals and $(I \setminus \{i\}, E')$ is the resulting tree after the algorithm removeLT(i) is performed. Consider intervals j and k in I, where $j \preceq k$. We want to show that $(j, k) \in E'$ if and only if (j, k) is an edge in $\mathcal{L}(I \setminus \{i\})$.

($\rightarrow$) Suppose $(j, k) \in E'$. We prove that (j, k) is an edge in $\mathcal{L}(I \setminus \{i\})$.

- Let $(j, k) \in E$. Assume that (j, k) is not an edge in $\mathcal{L}(I \setminus \{i\})$. By Lemma 10, either $j = i$ or $k = i$. If $j = i$, that is, i is a child of k in $\mathcal{L}(I)$, then the algorithm removes the edge (i, k) in line 2. Consider the case when $k = i$, that is, j is a child of i. If j is the equivalent child of i, the algorithm removes the edge (j, i) in lines 3-5. If j is the compatible child of i, the algorithm removes (j, i) in lines 10-12. In either case $(j, k) \notin E'$, which is a contradiction.

- Let $(j, k) \notin E$. The algorithm removeLT must have added the edge (j, k). There are four possible cases. First, the algorithm adds an edge in line 7, that is, $k = \mathsf{rc}(i)$. Then j is the equivalent child of i and i is the greatest interval in $[i]$. After removing i, j is the greatest interval in $[j]$, so that j is the compatible child of k. Second, the algorithm adds an edge in line 9, that is, k is the next interval of i. Then $i \sim k$. Since j is the equivalent child of i, $j \sim k$. Third, the algorithm adds an edge in line 17. Then k is the next interval of j with respect to $I \setminus \{i\}$ and $j \sim k$. Finally, the algorithm adds an edge in line 19. Then j is the greatest interval in $[j]$ and $k = \mathsf{rc}(j)$ with respect to $I \setminus \{i\}$. In all these case the edge (j, k) is an edge in $\mathcal{L}(I \setminus \{i\})$.

($\leftarrow$) Suppose (j, k) is an edge in $\mathcal{L}(I \setminus \{i\})$. We prove that $(j, k) \in E'$.

- Let $(j, k) \in E$. Assume that $(j, k) \notin E'$. Then the algorithm removeLT must have deleted the edge (j, k). First, the algorithm removes an edge from i (see line 2). Second, it removes an edge from the equivalent child of i (see lines 3-58). Finally, it removes an edge from the compatible child of i (see lines 10-12). Thus the algorithms removes only edges, incident to i, but these edges are not in $\mathcal{L}(I \setminus \{i\})$, which is a contradiction.

- Let $(j, k) \notin E$. By Lemma 10 (j, i) is an edge in $\mathcal{L}(I)$. Suppose j is the equivalent child of i. The algorithm finds j in lines 3-5. If j is the compatible child of k in $\mathcal{L}(I \setminus \{i\})$, then i is the compatible child of k in $\mathcal{L}(I)$. If j is the equivalent child of k, then k is the next interval of i. The algorithm takes care of both cases in lines 6-9 and adds the edge (j, k) in line 9. Suppose j is the compatible child of i. The algorithm finds j in lines 10-11. If j is the equivalent child of k, then k is the next interval of j and $\mathsf{rc}(j) = \mathsf{rc}(k)$ with respect to $I \setminus \{i\}$. The algorithm adds the edge

(j, k) in lines 15-17. If j is the compatible child of k, then the algorithm adds the edge in line 19. Thus, $(j, k) \in E'$. □

Lemmas 8-12 lead us to the following theorem:

Theorem 13. *The* queryLT, insertLT *and* removeLT *operations solve the dynamic monotonic interval scheduling problem in* $O(\log n)$ *amortised time, where* n *is the size of the set* I *of intervals.*

Note. The time complexity of the operations above depends on the type of dynamic trees, representing paths of LT. We can achieve the worst-case bound instead of amortized if we use globally biased trees instead of splay trees [11]. However, after each operation we must ensure that for every pair of edges (v, u) and (w, u) of the linearised tree, nodes v and u are in the same dynamic tree if and only if the numbers of nodes in the subtree rooter at v is greater or equal to the number of nodes in the subtree rooted at u.

5. Experimental results

In this section we present an experimental comparison between three algorithms for solving monotonic case of the dynamic interval scheduling problem: (i) the naive dynamic algorithm N, which keeps the intervals a binary search and calculate the greedy optimal set from scratch at every query operation; (ii) the algorithm CF based on the compatibility forest; (iii) the algorithm LT based on the linearised tree. We implemented these algorithms in Java. The algorithm N is based on the standard Java implementation of Red-Black tree, which we extended with left_compatible and right_compatible operations. We use the implementation of N in the algorithms CF and LT to store intervals and perform tree operations. In CF and LT we implemented bottom-up splay operation as described in [11]. We run the experiments on a laptop with *4GB of RAM* memory and *Intel Core 2 Duo 2130 Mhz, 3MB of L2 cache memory* processor.

In our experiments, we measure the total and the average running time of a sequence of m operations on initially empty interval set. The sequence consists of n insert operations, rn remove operations and qn query operation, where n is a linearly increasing number and r and q are fixed parameters of the experiment. We create a sequence of operations randomly while satisfying two conditions. First, whenever we invoke an insert operation of an interval i, we make sure that there is no interval i in the set. Second, whenever we invoke a remove operation of i, we make sure that i exists in the set. Thus every update operation calls for an actual change of the interval set.

To better understand the algorithms' performance, we defined the *sparsity* of an interval set I to be the upper bound on the ratio between the size of the greedy optimal set J and the size of I. The smaller the sparsity, the more intervals pairwise overlap. For example, if the sparsity is $1/2$, we make sure by creating intervals of the length $2/n$ that at most every second interval can belong to J.

The sparsity of I has an important influence on the algorithms N and CF. In the compatibility forest we conclude every update operation with the expose operation on the least interval in the set, which restores the missing edges between intervals from J. Therefore the smaller sparsity, the smaller chance of an update operation to affects the splay tree, representing set J. In the naive algorithm, the query operation may visit every interval from J. Therefore the smaller sparsity, the less maximal number of intervals the query operation may visit.

Experiment 1. The analysis of the algorithms shows that N updates the interval set faster than CF and LT, but queries the set slower. Therefore in the first experiment we measured the efficiency of the algorithms undergoing n insert, $0.5n$ remove and $0.01n$ query operations. The operations are shuffled as described above. We set the sparsity parameter to be 0.1. The result of the experiment is shown on the Figure 6.

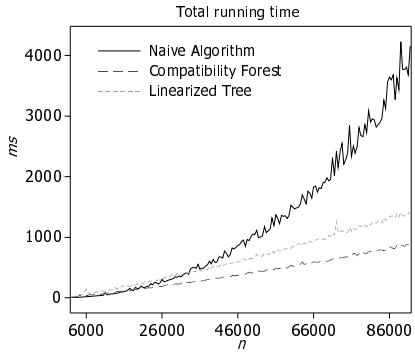

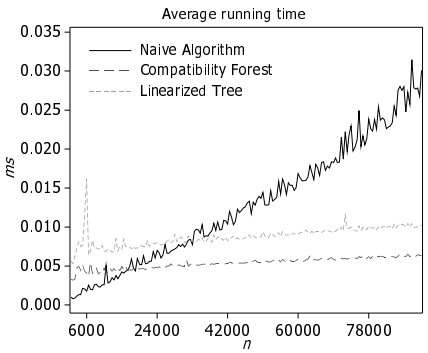

n	N	CF	LT
66000	1823	597	928
67000	1745	590	968
68000	1802	612	972
69000	1904	610	965
70000	1931	610	1027

n	N	CF	LT
66000	0.01829	0.00599	0.00931
71000	0.02155	0.00598	0.00998
76000	0.02032	0.00606	0.00956
81000	0.02371	0.00625	0.00981
86000	0.02803	0.00651	0.01027

Figure 6: sparsity is 0.1, $0.5n$ remove operations, $0.01n$ query operations, n insert operations.

The experiment shows that the difference of total running time between algorithms undergoing a sequence of operations with number of insertion less then 6000 is small, especially between CF and N. However, when we increase the number of insert operations, N performs much slower than two other algorithms. The average running time per operation of N is increasing similarly to a linear function, whereas the average running time per operation of CF and LT increases much slower. The experiment also shows that CF updates the interval set with low sparsity faster than LT.

Experiments 2 and 3. In the next two experiments we measure the performance of CF and LT undergoing a sequence of operations with the equal number of

insert and query operations. We excluded N from the experiments because N performs too slowly when the number of query operations increases. The difference between the second and the third experiment is in the number of remove operations. Sequences in Experiment 2 do not contain remove operations. Sequences in Experiment 3 contain $0.5n$ remove operations. We set the sparsity parameter to be 0.8. Figure 7 shows the results of Experiment 2, Figure 8 shows the results of Experiment 3.

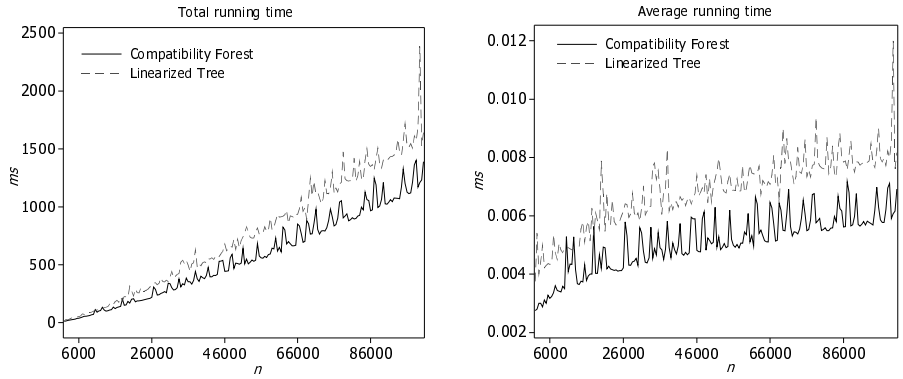

n	CF	LT
82000	901	1221
84000	975	1230
86000	966	1348
88000	995	1508
90000	1027	1412

n	CF	LT
82000	0.00549	0.00744
84000	0.00580	0.00732
86000	0.00561	0.00783
88000	0.00565	0.00856
90000	0.00570	0.00784

Figure 7: sparsity is 0.8, no remove operations, n query and insert operations.

The third experiment shows that if we do not allow remove operations, CF performs faster than LT. If we allow remove operations, CF performs slightly slower than LT. However, the results of Experiment 2 show that if the interval set is not sparse, CF inserts and removes intervals faster than LT.

Conclusion. The experimental result verifies our theoretical analysis and shows that both CF and LT runs significantly faster than the naive algorithm. Moreover, the results show that in a random environment CF performs as fast as LT to within a constant factor, despite the worst $(\log^2 n)$ time upper bound. Considering that CF is relatively easy to implement, CF can find its practical applications.

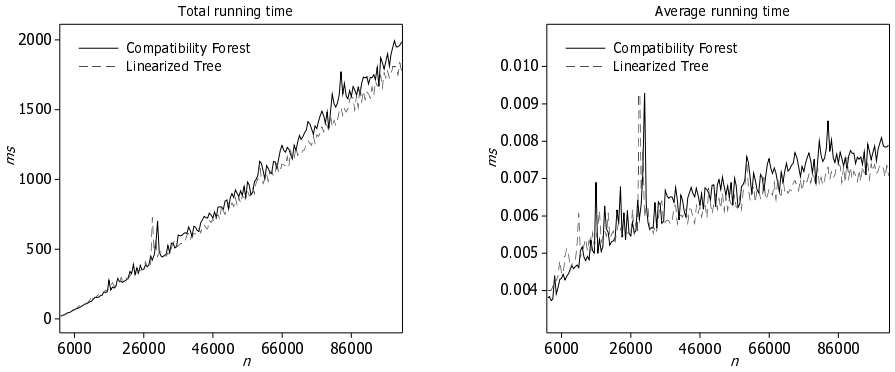

n	CF	LT
82000	1547	1425
84000	1688	1507
86000	1593	1587
88000	1660	1516
90000	1727	1572

n	CF	LT
82000	0.00754	0.00695
84000	0.00803	0.00717
86000	0.00740	0.00738
88000	0.00754	0.00689
90000	0.00767	0.00698

Figure 8: sparsity is 0.8, $0.5n$ remove operations, n query and insert operations.

References

[1] D. Corneil. A simple 3-sweep LBFS algorithm for the recognition of unit interval graphs. *Discrete Applied Mathematics*, 2004, 138(3), 371-379.

[2] X. Deng, P. Hell and J. Huang. Linear-time representation algorithms for proper circular-arc graphs and proper interval graphs. *SIAM Journal on Computing*, 1996, 25(2), 390-403.

[3] S. Fung, C. Poon, and F. Zheng. Online interval scheduling: Randomized and Multiprocessor cases. *Proceedings of COCOON*, 176-186, 2007.

[4] M.J. Katz, F. Nielsen and M. Segal. Maintenance of a piercing set for intervals with applications. *Algorithmica*, 36(1), 2003, 59-73

[5] R. Lipton and A. Tompkins. Online interval scheduling. *Proceedings of the Fifth Annual ACM-SIAM Symposium on Discrete Algorithms*, 1994, 302-311.

[6] P. Heggernes, D. Meister, and C. Papadopoulos. A new representation of proper interval graphs with an application to clique-width. *Electronic Notes in Discrete Mathematics*, 2009, 32, 27-34.

[7] H. Kaplan,E. Molad and R. Tarjan. Dynamic rectangular intersection with priorities. *Proceedings of the thirty-fifth annual ACM symposium on Theory of computing*, 2003, June, 639-648.

[8] J. Kleinberg and E. Tardos. *Algorithm Design*, 2006.

[9] A. Kolen, J.K. Lenstra, C.H. Papadimitriou, and F.C. Spieksma. Interval scheduling: A survey, *Naval Research Logistics*, 54, 5, 2007, 530–543.

[10] Mehlhorn, K. *Data structures and algorithms, Volume 3: Multi-dimensional Searching and Computational Geometry.* Springer-Verlag, Berlin, 1984.

[11] D. Sleator and R. Tarjan. A Data Structure for Dynamic Trees, *Journal of computer and system sciences*, 26, 3, 1983, 362–391.

[12] D. Sleator and R. Tarjan. Self-adjusting binary search trees, *Journal of the ACM*, 32, 3, 1985, 652–686.